Eco514—Game Theory
1. Game Theory = Multiperson Decision Theory; Zero-Sum games

Marciano Siniscalchi

September 16, 1999

Administrative Stuff

Class: Tue-Thu 10:40-12:10 [?], Room 317, Bendheim. OH, by appointment.

The Big Picture

Most of you will already have used some of the tools of GT in your core courses. You will probably be familiar with the notions of *simultaneous vs. extensive-form game*, *perfect vs. imperfect information*, *complete vs. incomplete information*, *Nash equilibrium*, *Perfect Bayesian equilibrium*, *Sequential equilibrium*, etc. (but, don't worry if you are not).

This course has two main objectives: (1) introduce additional, perhaps more sophisticated analytical tools, which you will most likely find useful in your own work; (2) enhance your understanding of the main *ideas* (as opposed to the main *tools*) of GT, which includes (among other things) a critical appraisal of the *applicability* of the techniques you have learned.

Under point (1), I will discuss, among other issues: sophisticated models of *incomplete information games* which, for instance, allow one to think clearly about currency speculation, bank runs and similar phenomena having to do with "market sentiment"; *forward induction*, a powerful idea which pops up in many applied models, from political economy to industrial organization; *non-equilibrium solution concepts*, which we may view as a way to test the robustness of strategic predictions to deviations from "textbook" settings; *reputation phenomena*, which may shed some light on, say, a certain software behemoth's unduly response to new entrants; and so on.

VISIT…

However, in my mind, (2) is much more important than (1). The reason is that, in my perhaps-not-so-humble opinion, GT is best viewed not as a "bag of tricks", but as a *methodology*—a select collection of internally consistent, general principles which helps you structure your *own* reasoning about strategic interaction in a given setting.

Games as Multiperson Decision Problems

Decision theory enjoys an enviable status in economics. It is, often seen as the prototypical example of a rigorous, well-founded theory, which is easy to apply to the analysis of a wide variety of relevant decision situations in the social sciences: from Econ 101-type models to GE and finance.

The link with GT might not be apparent. Indeed, it has been somewhat downplayed in the past, because certain conclusions one can draw from strategic reasoning seem to be so much at odds with decision theory that the very existence of such a link might appear questionable (think about: "more information can hurt you.")

However, the link is definitely out there, and recent research has brought it to the fore. What might be surprising is that the link was quite evident in the early days of GT, when people were mostly interested in *zero-sum games*. I shall presently state my case—and, in the process, review a related mathematical tool whose importance I simply cannot overemphasize: linear programming.

Making the connection

First, I need to define a *decision problem*.

Definition 1 *A decision problem is a tuple* $\{\Omega, C, F\}$*, where* Ω *is a set of possible* states of the world, C *is a set of* consequences*, and* $F \subset C^{\Omega}$ *is a set of* acts.

The idea is that, if the decision-maker (DM) chooses act $f \in F$, and the prevailing state is $\omega \in \Omega$, then the consequence $f(\omega) \in C$ obtains. Consequences are supposed to represent all the DM cares about.

The DM does not observe the prevailing state of the world: that is, there is *uncertainty* about it. Still, she might be assumed to exhibit preferences (represented by a binary relation $\succeq \subset F \times F$) among available acts. The classical "representation" result exhibits a set of joint assumptions about $\succeq$ (and F) which are equivalent to the existence of a probability measure $\mu \in \Delta(\Omega)$ and a utility function $u : C \to \mathbf{R}$ (unique up to a positive affine transformation) which represents $\succeq$:

$$\forall f, g \in F, \quad f \succeq g \Leftrightarrow \int_{\Omega} u(f(\omega))\mu(d\omega) \geq \int_{\Omega} u(g(\omega))\mu(d\omega)$$

Next, let us define *simultaneous games*:

Definition 2 *A finite simultaneous game is a tuple* $\{N, (A_i, u_i)_{i \in N}\}$*, where* N *is a finite set of players, and for each player* $i \in N$*,* A_i *is a finite set of actions and* u_i *is a payoff function* $u_i : A_i \times A_{-i} \to \mathbf{R}$*.*

(I use the conventional notation $A_{-i} = \prod_{j \neq i} A_j$ and $A = \prod_{i \in N} A_i$.) We could be more general and assume that each player is endowed with a preference relation over action profiles, but let's not worry about this right now.

The above is the "standard" definition of a game. As you may see, it is not quite consistent with the definition of a decision problem, so we need to reinterpret things a bit.

Indeed, the most important difference is that a game implicitly defines not just one, but N decision problems (with a conventional abuse of notation, I will denote by this both the *set* N and its cardinality $|N|$).

Let's take the point of view of Player i. She has no control over her opponents' choices, and she will not find out what they were until the game is over. Thus, A_{-i} seems a good candidate for the state space: $\Omega = A_{-i}$.

Next, utilities are attached to points $(a_i, a_{-i}) \in A$ according to Definition 2. This amounts to saying that *there is a one-to-one correspondence between the actual outcome of the game* (which is what players care about) *and action profiles*. Thus, we may as well imagine that action profiles are the relevant consequences: $C = A$.

By a process of elimination, acts must be actions: "$F = A_i''$. Formally, action $a_i \in A_i$ defines an act $f_{a_i} : A_{-i} \to A_i \times A_{-i}$ (i.e. $f_{a_i} : \Omega \to C$) by $f_{a_i}(a_{-i}) = (a_i, a_{-i})$.

However, the definition of a game incorporates additional information: specifically, it includes a utility function for each player. This implies that we have a complete description of players' preferences among *consequences*. However, in view of the representation results cited above, this is only one part of the description of players' preferences among *acts*, because *the probabilities are missing*.

This trivial observation is actually crucial: in some sense, the traditional role of game theory has been that of *specifying those probabilities* in a manner consistent with *assumptions or intuitions concerning strategic rationality*.

Nash Equilibrium

To make this more concrete, recall what you may have heard about Nash equilibrium: a player's strategy doubles as her opponents' beliefs about her strategy.

Thus, in a 2-player game, a Nash equilibrium is a specification of beliefs (α_1, α_2) for both players, with the property that if Player i's belief assigns positive probability to an action $a_j \in A_j$ $(j \neq i)$, then the act f_{a_j} is preferred by j to any other act, given that j's preferences are represented by the utility function $u_j : A \to \mathbf{R}$ and the

probability measure $\alpha_i \in \Delta(A_i)$ (note the indices!). In short, we say "a_j is a best reply to α_i".

Note: at first sight, Nash equilibrium seems to say that players' beliefs are required to be "consistent with rationality". However, note that, since the rationality of an act can only be assessed with respect to a given (utility function and) probability measure over states of the world, the definition really says that players' beliefs are required to be "consistent with rationality, *given the equilibrium beliefs*." We shall be more specific about this point in a subsequent lecture.

Also, note that this interpretation of Nash equilibrium has *no behavioral content*: it is simply a restriction on *beliefs*. If we really wish to push this to the extreme, we can say that, if players' belief satisfies the restrictions in the definition of a Nash equilibrium, then the players may choose an action that is a best reply to those beliefs. Again, we will return to this point later in the course.

Zero-Sum games and the minmax solution

Game theorists were originally interested in a much simpler class of games, namely two-player games in which interests were diametrically opposed:

Definition 3 *A two-player simultaneous game* $\{\{1,2\},(A_i,u_i)_{i\in N}\}$ *is zero-sum iff, for all* $a \in A$, $u_1(a) + u_2(a) = 0$.

To put this in perspective, think about board games such as chess or checkers, or war.

In the latter case, it appears that *doing one's best in the worst-case scenario* might be the "right" thing to do. I am told this is what they teach you in military academy, anyway.

In a zero-sum game, the worst-case scenario for Player 1 is also the best scenario for Player 2, and vice-versa. Thus, for Player 1 to assume that Player 2 is "out to get her" is not inconceivable—after all, Player 2 has all the right incentives to do so.

This idea leads to the notion of *maxmin* and *minmax* strategies. Suppose first that both players can *randomize*—commit to spinning roulette wheels or something like that, so as to choose strategies according to any pre-specified probability $\alpha_i \in \Delta(A_i)$. We will get back to this later in the course: for now, let's just assume that this is possible; the reason why this is convenient will become clear momentarily. Again using a conventional notation, let $u_1(\alpha_1,\alpha_2) =: \sum_{(a_1,a_2)\in A} u_1(a_1,a_2)\alpha_1^1(a_1)\alpha_2^1(a_2)$ for any $(\alpha_1,\alpha_2) \in \Delta(A_1)\times\Delta(A_2)$.

Consider first a pair of mixed strategies $\alpha_i^1 \in \Delta(A_i)$, $i = 1,2$, such that

$$u_1(\alpha_1^1,\alpha_2^1) = \min_{\alpha_2\in\Delta(A_2)} \max_{\alpha_1\in\Delta(A_1)} u_1(\alpha_1,\alpha_2)$$

In words, $u_1(\alpha_1^1, \alpha_2^1)$ is the minimum payoff that Player 1 can secure; equivalently, it is the minimum payoff to Player 1 when Player 2 attempts to damage her as much as possible, while keeping into account the fact that Player 1 best-responds.

Note that since $\max -f = -\min f$,

$$u_2(\alpha_1^1, \alpha_2^1) = -u_1(\alpha_1^1, \alpha_2^1) = \max_{\alpha_2 \in \Delta(A_2)} \min_{\alpha_1 \in \Delta(A_1)} u_2(\alpha_1, \alpha_2)$$

i.e. (α_1^1, α_2^1) also yields Player 2 the maximum payoff he can obtain if Player 1 chooses her mixed strategy so as to hold Player 2's payoff down.

Such a pair of mixed strategies can always be found (why?). Similarly, one can consider the symmetric problem for Player 2, and define a new pair (α_1^2, α_2^2) such that

$$u_2(\alpha_1^2, \alpha_2^2) = \min_{\alpha_1 \in \Delta(A_1)} \max_{\alpha_2 \in \Delta(A_2)} u_2(\alpha_1, \alpha_2)$$

so $u_2(\alpha_1^2, \alpha_2^2)$ is the minimum payoff that Player 2 can secure; as above,

$$u_1(\alpha_1^2, \alpha_2^2) = \max_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in \Delta(A_2)} u_1(\alpha_1, \alpha_2)$$

the maximum payoff that Player 1 can obtain if Player 2 attempts to hold Player 1's payoff down.

It may be helpful to think about this as a "pseudo-dynamic" story: in order to define (α_1^1, α_2^1), one imagines that Player 2 chooses first, and Player 1 moves next; the latter best-responds to Player 2's choice, and the former, anticipating 1's best response, attempts to minimize Player 1's payoff. On the other hand, to define (α_1^2, α_2^2), one imagines that Player 1 chooses first, and Player 2 moves next; the latter observes 1's choice and attempts to minimize 1's payoff, while the former, anticipating this, attempts to maximize her own payoff. [Of course there is a dual to this story, with 1's and 2's roles inverted.]

It is natural to ask whether Player 1's resulting payoff under the two definitions is the same. The answer is affirmative:

Theorem 0.1 *(a.k.a. "The minmax theorem"): there exists a pair* $(\alpha_1^*, \alpha_2^*) \in \Delta(A_1) \times \Delta(A_2)$ *such that*

$$u_1(\alpha_1^*, \alpha_2^*) = \max_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in \Delta(A_2)} u_1(\alpha_1, \alpha_2) = \min_{\alpha_2 \in \Delta(A_2)} \max_{\alpha_1 \in \Delta(A_1)} u_1(\alpha_1, \alpha_2)$$

The easiest proof of this theorem uses Linear Programming. Let n_i be the cardinality of A_i and number its elements from 1 to n_i; define the $n_1 \times n_2$ matrix U by letting $U_{k\ell} = u_1(a_1^k, a_2^\ell)$, where a_1^k and a_2^ℓ are the k-th and ℓ-th elements of A_1 and A_2, respectively. Then Player 1's (2's) beliefs about Player 2's choices may be represented as nonegative vectors in $\alpha_2 \in \mathbf{R}^{n_2}$ (resp. $\alpha_1 \in \mathbf{R}^{n_1}$) whose elements add up to one.

Consider the following problem:

$$\min_{\alpha_2} u \quad \text{s.to} \quad -U \cdot \alpha_2 + 1u \geq 0,\ 1'\alpha_2 = 1$$

which we can interpret as follows: Player 2 chooses α_2 so that every strategy of Player 1 yields u or less; indeed, Player 2 attempts to make u as small as possible. Note that u is unconstrained in sign.

The dual program is

$$\max_{\alpha_1} v \quad \text{s.to} \quad -\alpha_1' U + 1'v \leq 0,\ \alpha_1' 1 = 1$$

i.e. Player 1 chooses α_1 so each of Player 2's strategy yields at least v to Player 1 (think about it!), i.e. at most $-v$ to Player 2. Again, v is unconstrained.

Note that both programs are feasible (easy); also, both programs are bounded, because, say, in the dual program, v cannot be larger than the largest entry in the matrix U (which, if you think about it, is also obvious!) But then, by the Existence-Duality theorem of LP, both programs have optimal solutions, the optimal values are the same (so $u = v$), and the optimal solutions are complementary. In particular, this means that $\alpha_1(a_1) > 0$ implies $u_1(a_1, \alpha_2) = u$, and similarly for 2. Thus, only actions which achieve the maximal payoff u against α_2 receive positive probability; in other words, α_1 only assigns positive probabilities to best responses (and similarly for α_2). Hence, the optimal solutions of the dual LPs identify a Nash equilibrium; it is then easy to see that any Nash equilibrium satisfies the equality in Theorem 1 (see Proposition 22.2 in OR if you don't see this), and we are done.